

ABSTRACT

In some embodiments, a method and apparatus for an automatic thread-partition compiler are described. In one embodiment, the method includes the transformation of a sequential application program into a plurality of application program threads. Once partitioned, the plurality of application program threads are concurrently executed as respective threads of a multi-threaded architecture. Hence, a performance improvement of the parallel multi-threaded architecture is achieved by hiding memory access latency through or by overlapping memory access with computations or with other memory accesses. Other embodiments are described and claimed.